



LynxOS Hot Swap for MCP750 PPC Systems Driver

with Embedded RAINlink Technology

Installation and User's Guide

Overview

This ZNYX Networks driver is provided for ZNYX NetBlaster™ Hot Swap Ethernet Adapters. It requires the production release of LynxOS 3.1.0A-HAP for PPC, or LynxOS 3.0.1 for PPC (release 110298-G) together with the evaluation version of the LynxOS 3.0.1HS SSP.

The appropriate driver can be obtained from the ZNYX web site: <http://www.znyx.com>. The distribution is available as a compressed, tar file. The driver includes embedded RAINlink technology. RAINlink provides critical link services such as link aggregation (trunking), link failover and dynamic load balancing over multiple network ports. These ports can be a group of multiple single-channel or multi-channel ZNYX adapter ports, up to a maximum of 16 ports.

The ZNYX LynxOS Hot Swap driver is dynamically loaded into the LynxOS kernel when supported ZNYX devices are found by the Hot Swap Event Manager (hsem). The following steps will take you through the additions necessary to make this driver available.

Systems Requirements

- Platforms:** Motorola MCP 750 based system with hot-swap ready CompactPCI slots (Motorola CPX2208, CPX8216)3
- Operating Systems:** LynxOS 3.1.0A-HAP for PPC
LynxOS 3.0.1 for PPC (release 110298-G) together with the evaluation version of the LynxOS 3.0.1HS SSP

Supported ZNYX NetBlaster Adapters:

Model	Bus	Number of Ports
ZX212 (Requires Motorola CPV8540 Carrier)	PMC	2
ZX214 (Requires Motorola CPV8540 Carrier)	PMC	4
ZX222(Requires Motorola CPV8540 Carrier)	PMC	2
ZX244 (Requires Motorola CPV8540 Carrier)	PMC	4
ZX442 hot-swap	cPCI	2

ZX444 hot-swap	cPCI	4
ZX474 rear i/o & hot-swap	cPCI	4
ZX478 rear i/o & hot-swap	cPCI	8

Retrieving the ZNYX Drivers

1. Go to the ZNYX web site at www.znyx.com and navigate under the driver downloads area to get the LynxOS driver for your adapter. Be sure to choose the hot-swap driver for the MCP750.
2. The driver is available as a compressed tar file. Retrieve the file to the `/src` directory of your LynxOS system. *Uncompress* and *untar* the file. You must have super-user privileges to complete this installation.

```
cd /src
tar -xvpzf ZNYXnb-hs.PowerPC.tar.Z
```

Installing the ZNYX Driver for LynxOS

1. If you have a previous version of the ZNYX driver for LynxOS installed, you must first uninstall it. The most important step is to remove or comment out the line "I:zxe.cfg" from the file `/sys/lynx.os/CONFIG.TBL`. You must also have installed the Lynx Hot Swap I/O Manager (HASW 1.0).
2. The driver installation files will be extracted and placed into a subdirectory called `ZNYXnb`. Run the program `Install.znb` from that directory:

```
cd ZNYXnb
./Install.znb-hs
```

This program will copy the driver files into their appropriate locations, set up a directory `/etc/rain`, for interface configuration information, and describe the remaining steps you must perform. The steps are detailed here.

- a. For installation in a Motorola CPX2208 System:

The original release of HASW 1.0 contains files designed for the Motorola CPX8216 systems, which provide 1 or 2 CPUs with Hot Swap Controllers (HSC). To use this release in a Motorola CPX2208 system, which has a single CPU and no Hot Swap Controller, several files must be replaced. These include the hot swap utility routines:

```
/usr/sbin/hasw/dr
/usr/sbin/hasw/hsem
/bin/hsls
```

and the configuration files:

```
/etc/hasw/hsem.conf
/etc/hasw/ha_sim.conf
```

The CPX2208 versions of these files are in:

```
/etc/rain/cpx2208
```

They can be installed by running:

```
/etc/rain/cpx2208/doInstall
```

The old files will be saved in `/etc/rain/cpx8216`.

- b. For installation in a Motorola CPX8216 System:

The file `/etc/hasw/hsem.conf` must be edited to add the hot swap configuration information for ZNYX Ethernet Adapters. The lines which must be added are in the file:

```
/etc/rain/zxe_hsconf.add
```

Edit the file `/etc/hasw/hsem.conf` and insert this file immediately after the line:

```
drivers
```

3. When multiple ethernet interfaces are added, additional system buffer resources will be needed to allow maximum performance. Edit the file `/sys/devices/hbtcpip_info.c` to adjust this resource allocation by making the following changes:
 - a. Increase the number of MBUFFs to 1024 if 8 or more interfaces are installed.
 - b. Increase the number of CLUSTERS to between 1/4 and 1/2 the MBUF value.
 - c. For better performance when RAINLINK trunking is used, the values for `tcp_sendspace` and `tcp_recvspace` should be increased to 49152.

After making these changes, save the file and rebuild the devices library:

```
cd /sys/devices  
make
```

4. If you made changes to the `hbtcpip` files in the previous step, you must build a new LynxOS kernel. Change directories to the kernel build directory and build a new kernel:

```
cd /sys/lynx.os  
make install
```

5. Next, boot the new kernel. This is necessary only if a new kernel was built in the previous step, since Znyx Hot Swap Adapters can be installed dynamically in any empty cPCI slot while the system is running. Any ZX404 or ZX408 rear panel transition cards must be installed before the adapter.

To install the adapters at this point, shutdown the system with the command

```
reboot -haN
```

Turn the power off when the shutdown completes and install the adapters and transition cards in available slots. Then restart the system. If the adapters are already installed, or will be installed later, and the system must be rebooted, use the command:

```
reboot -aN
```

When the Hot Swap Event Manager is started (by `ha_sim.conf`), any Hot Swap adapters are discovered and the device install program defined for the adapter in `/etc/hasw/hsem.conf` is invoked. For Znyx adapters, the `zxe_dev_install` program will load the Znyx device driver, if it is not already loaded, and add an entry to the `/dev` directory, `/dev/zrm`, for use in communicating with the driver from applications such as `rlconfig` and `zxe_hscctl`. Once the driver is loaded, each new adapter instance is installed.

By default, each twisted pair network interfaces will auto-negotiate its link speed, and full or half-duplex operation. If a fixed setting is desired, edit the media variable in the `zxe_dev_install` program to specify the desired setting. The possible settings are listed in the program. This program is located in the directory:

```
/sys/drivers.cpci_drm/if_zxe
```

As a Znyx adapter is installed, a network interface is created for each port on the adapter. The list of currently available interfaces is given by the command:

```
netstat -i
```

The Znyx adapter entries are `znb0`, `znb1`, etc.

6. You will need to supply an IP address or name for each interface. See your network administrator for details about appropriate names. Add these entries to the `/etc/hosts` file.
7. To define and configure RAINlink interfaces, run the program:

```
/etc/rain/rainlink -a
```

Basic RAINlink configurations are described in the next section. All network interfaces which are to be part of the Rainlink configuration must be present when this program is run. If you are not using any rainlink features, run the program:

```
/etc/rain/rainlink -d
```

8. A utility routine, `/etc/rain/showports`, can be used to show which ZNYX adapters are present in the system, and what interface names have been assigned to each port on the card. A `znb` interface name that is not in a trunk or failover group is available for configuration. The rainlink trunks and failover groups are configured using their interface names, which begin with `zrl`.
9. The startup file, `/etc/rain/rainrc`, can be invoked at system startup to configure the Znyx adapter interfaces which are present.

First edit the file `/bin/rc` to move the invocation of `/usr/sbin/hasw/ha_sim` above the call to `/net/rc.network`, so that the Hot Swap Event Manager will run first, allowing the

adapters which are present to be discovered and the network interfaces created. Then, edit the file `/net/rc.network` to add the line:

```
startit /etc/rain/rainrc
```

in the network interfaces section.

If any of network interfaces included in the `rainrc` file are not be present at system startup, do not add the "startit" line to `/net/rc.network`. The command

```
/etc/rain/rainrc
```

can be run later, after the adapters are installed.

10. Znyx Hot Swap capable adapters can be added when the system is running, by inserting the card into an available slot. If an adapter using a rear panel transition card is installed, the rear panel transition card must be installed first, before the adapter is installed. The insertion of the adapter should be recognized by the system as soon as the lower ejector handle is fully raised and locked. The Hot Swap Event Manager will run the `zxe_dev_install` program as described above. Once the adapter is installed, its ports can be configured.
11. When a Znyx Hot Swap adapter is to be removed, the operator must first unlock the lower ejector handle. When the system detects the change, the Hot Swap Event Manager will run the `zxe_dev_uninstall` program. The network ports on the adapter will be declared down, causing any Rainlink configurations to fail over to standby interfaces; the PCI device resources allocated to the ports and the adapter will be released, and the blue light on the adapter turned on. The adapter can then be removed completely by pushing the top and bottom ejector handles up and down respectively.
12. The network interfaces associated with ports on the adapter are not removed, and remain configured if they were before the adapter was removed; the link state will be down, just as if the cable were removed. If a replacement adapter of the same type is inserted in the slot, the `zxe_dev_install` program will not install a new device instance, but instead signal the driver to reallocate PCI device resources for the new card and attempt to bring the links back up, ready to communicate through the already existing network interfaces.
13. To completely remove an adapter from the system, deleting its network interfaces, a utility program, `zxe_clrslot`, is provided. The adapter itself should be removed from the slot, then run:

```
/etc/rain/zxe_clrslot <s>
```

where `<s>` is the slot number of the adapter to be removed. All `znb` and `zrl` interfaces involving any devices in the slot are configured down, any routes using the interfaces are removed, and rainlink configurations deleted before device removal. The slot and network interface names are then available for reuse.

14. To unload the Znyx driver, use the program `zxe_unload`.

```
/etc/rain/zxe_unload
```

This will also clear all Znyx adapter entries, so it is not necessary to first use `zxe_clrslot`.

Installing RAINlink for LynxOS

All the necessary installation and configuration programs for RAINlink for LynxOS are placed on the system during the installation of the NetBlaster driver.

1. RAINlink for LynxOS includes utilities and example configuration files for standard configurations in order to get you up and running quickly. The necessary files are placed onto your system during the NetBlaster driver installation in */etc/rain*. Change directories to */etc/rain* to begin the installation:

```
cd /etc/rain
```

2. An installation script is included with RAINlink for LynxOS, called *rainlink*. The *rainlink* script can be used to activate, deactivate, or display the current RAINlink for LynxOS configuration. The first time the *rainlink* script is run, the script will give you a list of default configurations to choose from suitable for most basic configurations of two or four port trunks or failover groups. Detailed configuration instructions are included in the next section. Start the *rainlink* script by typing *rainlink -a* at the command prompt:

```
rainlink -a
```

3. The script dialogue will allow you to choose from a set of default configurations.

RAINlink

This script initializes or reinitializes the RAINlink layer within the ZNYX NetBlaster driver. An input script located at */etc/rain/rainlink.conf* is used as input to ZNYX's *rlconfig* application. Either choose one of the following prepared input scripts which will be copied to */etc/rain/rainlink.conf*, or exit and create your own custom script.

- 1 - System-to-System trunking between 4 ports.
- 2 - Fast EtherChannel trunking between 4 ports.
- 3 - Fast Failover between 4 ports.
- 4 - System-to-System trunking between 2 ports.
- 5 - Fast EtherChannel trunking between 2 ports.
- 6 - Fast Failover between 2 ports.
- 7 - Exit and create custom script.

Choose one of the above 7 choices(1-7):

4. Make your selection and continue. A note will be displayed informing you that the *rainlink* script will create a configuration file named, */etc/rain/rainlink.conf*, which will be the default configuration file used from this point forward. Details of the contents of the configuration file are provided later. Hit Enter to continue:

```
NOTE: /etc/rain/rainlink.conf has been created. If a rainlink.conf file
exists rainlink will not prompt for one of the previous 7 choices
To modify the RAINlink configuration in the future, edit
/etc/rain/rainlink.conf, and either run rainlink again, or reboot the
system.
```

```
Hit any key to continue:
Bringing down all Net Blaster interfaces now....
Running rconfig ...
#####TRUNK0 Successfully Created
TRUNK0 mode set
```

5. The next series of questions relate to bringing up the interfaces. Next, bring up the RAINlink interfaces, and then the remaining non-RAIN interfaces if there are any. You will need to supply IP names or addresses for each configured interface you bring up.

```
Do you want to configure and bring up the RAINlink interfaces [y,n,?] y
Enter hostname for interface zrl0 [server-zrl0]: server1
```

```
The following host names must be in the hosts database.
Check with your Network Administrator for more information.
```

```
server1
```

```
Do you want to configure and bring up any non-RAINlink interfaces [y,n,?] y
```

```
Bring znb4 up [y,n,?] y
Enter hostname for interface znb4 [server-znb4]: access1
```

```
The following host names must be in the hosts database.
Check with your Network Administrator for more information.
```

```
access1
```

```
Bringing up available boards now....
```

The script then exits. Your RAINlink and non-RAINlink interfaces should now be configured.

6. You can use *rainlink -s* to verify the status of your RAINlink configuration:

```
rainlink -s
```

```
trunk0<LAYER_3_MODE> = znb0<UP>, znb1<UP>, znb2<UP>,  
znb3<UP ==> zrl0
```

Configuring RAINlink for LynxOS

Typically you need to configure either trunks or failover groups for a given network. With RAINlink for LynxOS, you can also include trunks in failover groups. It is strongly recommended that you use trunking whenever your network adapters connect directly system-to-system, or to Cisco FEC compatible switch or router or other adapters. This is because trunking provides both increased bandwidth and improved fault-tolerance, whereas the failover service only provides improved fault-tolerance.

Configuring a trunk means to create (add) a trunk between two end stations and assign the member ports (links) to a trunk from EACH end station. A trunk should have at least two member ports (links.) Each port in a trunk can be a half or full duplex.

Configuring a failover group means to create (add) a failover group between two end stations and assign the member ports (links) to a group from each station. Either ports or trunks, or both can be members of a failover group. A failover group should have at least two members. Each port in a failover group can be a half or full duplex.

Quick Start

For most configurations, you will only need to pick the correct default configuration profile and run *rainlink -a*, as described previously.

If you wish to customize configuration for trunks or choose a different scheme for the dynamic load balancing, go to the "Advanced Trunking" section. If the link aggregation service is not suitable to your network configuration, then go to the "Advanced Failover" section to configure the failover groups instead of trunks. You can also configure failover groups of trunks themselves. Always configure trunks before failover groups.

Advanced Trunking

You can create link aggregation groups (trunks), assign ports to these trunks and select a scheme of dynamic load balancing. The most straightforward method for configuring RAINlink for LynxOS is to edit the */etc/rain/rainlink.conf* file and run *rainlink -a*. The script de-installs the current ZNYX interfaces, runs a configuration application *rlconfig* with the *rainlink.conf* input script, and brings up the RAIN links. The *rainlink -a* script can also bring up any regular non-RAIN interfaces, which were not configured into the RAIN links. All these steps can be done individually. Details of using the *rlconfig* application are included in a later section.

To add a trunk, edit the file */etc/rain/rainlink.conf*. The file is a plain text file that is extensively commented. Commented lines begin with a pound sign (#). Add a line for each trunk containing the list of ports to be included in that trunk. End each line with a semicolon. For example, to include ports 0 through 3 in trunk0, add an entry:

```
trunk0=znb0, znb1, znb2, znb3;
```

The order of entries is not important. You can use any available ports in any order. Port entries should not be duplicated in trunks. A trunk must consist of at least two ports, and cannot contain other trunks. The following is an acceptable way to configure two trunks:

```
trunk0=znb0, znb3;  
trunk1=znb1, znb2;
```

For system-to-system trunking, enable IP Trunking mode. Both systems should be setup identically. In IP Trunking mode, packets are sent to the driver as large datagrams and fragmented across the active links. To enable IP Trunking mode, uncomment/add the following line for each configured trunk:

```
trunk0=ip_trunking_mode;
```

For system-to-switch trunking, the default method of load balancing scheme uses Layer 3 Protocols. To use only the “Layer 2 Protocols” for dynamic load balancing, comment out the “layer_3_mode” option, and uncomment the “layer_2_mode” option. The options are mutually exclusive.

```
# trunk0=layer_3_mode;  
trunk0=layer_2_mode;
```

To explicitly choose Layer 3, uncomment (or add) the “Layer 3 Protocols” option, and comment out the “Layer 2 Protocols” option:

```
trunk0=layer_3_mode;  
# trunk0=layer_2_mode;
```

For system-to-switch trunking, you can disable Balance mode. Balance mode is on by default, and is normally left enabled. Balance mode dynamically distributes the load across the ports equally. It can be disabled by adding a “not equals” to the option:

```
trunk0 != balance_mode;
```

An “equals” sign explicitly enables balance mode:

```
trunk0 = balance_mode;
```

Once you have assembled your *rainlink.conf* file, run *rainlink -a* to bring up RAINlink for LynxOS with the new configuration.

```
rainlink -a
```

Advanced Failover

You can create link failover groups, assign ports and trunks to these groups and select a modes of failover. Failover groups are configured in the same manner as trunks: Edit the */etc/rain/rainlink.conf* file and run *rainlink -a*. *Rainlink -a* de-installs the current ZNYX interfaces, runs a configuration application *rlconfig* with the *rainlink.conf* input script, and brings up the RAIN links. The *rainlink -a* script can also bring up any regular non-RAIN interfaces, which were not configured into the RAIN links. All these steps can be done individually. Details of using the *rlconfig* application are included in a later section.

To add a failover group, edit the file `/etc/rain/rainlink.conf`. The file is a plain text file that is extensively commented. Commented lines begin with a pound sign (#). Failover groups can contain individual ports, trunks, or both ports and trunks. Always configure trunks before failover groups. Add a line for each failover group containing the list of ports and trunks to be included in that group. End each line with a semicolon. For example, to build a failover group of ports 0 and 1:

```
failover0 = znb0, znb1;
```

To build a failover group of two trunks, first build the trunks, then the failover group:

```
trunk0 = znb0, znb1;
trunk1 = znb2, znb3;
failover0 = trunk0, trunk1;
```

The following order will not work. You must build the trunks first:

```
failover0 = trunk0, trunk1;
trunk0 = znb0, znb1;
trunk1 = znb2, znb3;
```

You can also mix trunks and ports in failover groups, as long as the trunks are built first:

```
trunk0 = znb0, znb1;
failover0 = trunk0, znb2, znb3;
```

To remove a failover group, simply remove it, or comment it out, and run `rainlink -a`.

The default mode of failover is fast failover. In this mode, RAINlink for LynxOS moves the traffic over to a redundant stand-by link in case of a link failure in as little as 500 milliseconds. In addition to fast failover, Timeout mode can be enabled for a failover group. By enabling Timeout Mode, if no traffic is received in the specified time interval, the active port is automatically switched to another available link. You can choose different failover modes for different groups. You can also specify the timeout interval. To set Timeout Mode on a failover group with a timeout of 30 seconds, add a line for the failover group:

```
failover0 = timeout_mode 30000;
```

Using Rlconfig

The `rlconfig` application views or changes the RAINlink for LynxOS configuration. `Rlconfig` is installed into `/etc/rain` during the NetBlaster driver installation. You can view the current RAINlink for LynxOS configuration at any time, but the interfaces must be “down” in order to change the configuration. See *ifconfig(1M)* for more explanation of how to bring an interface down.

Options for `rlconfig` include:

Option:	Use:
-s	Displays the current RAIN configuration. Commands are not read from standard input with this option.

-t	Tear down all previously configured RAIN ports. Commands are not read from standard input with this option.
-l [#] file_name	Creates a file containing a list of the configured ZRL devices and the available ZNB devices. The “#” is used for version identification. This option is not intended for use by the user; it is used by the rainlink script to configure RAIN
-p ppa	Displays hardware level statistics for specified ppa device.

To view the current configuration at any time, enter:

```
rlconfig -s
```

The resulting display shows you the status of the RAIN configuration.

```
trunk0<IP_TRUNKING> = znb0<UP>, znb1<UP>, znb2<UP>,
znb3<UP ==> zrl0
```

Lines can be input directly to *rlconfig* from standard input, or from a script like */etc/rain/rainlink.conf*.

To start *rlconfig*, enter:

```
rlconfig
```

Enter commands one line at a time and conclude with Ctrl-d.

To use a previously prepared file of commands, redirect the file into *rlconfig*. For example:

```
rlconfig < /etc/rain/rainlink.conf
```

Typical usage would be to redirect a script of commands into *rlconfig* which instruct *rlconfig* to build and set modes on trunks and failover groups. Semicolons delimit commands. Spaces and new lines are ignored. Commands take the following form:

```
trunk<number> = znb<number> [ , znb<number> ... ] ;
failover<number> = znb | trunk<number> [ , znb|trunk<number> ...]

trunk<number> = mode ;
trunk<number> != mode ;

failover<number> = mode ;
failover<number> != mode ;
```

where <number> is a value between 0 and the maximum number of ports in the system. Trunks can consist of ports denoted by a comma-delimited list of znb<number>. A failover group can consist of ports or trunks. You must build a trunk prior to its use in a failover group.

Acceptable modes for trunks are:

Mode:	Appropriate for:
ip_trunking_mode	System-to-system trunking

layer_2_mode	System-to-switch trunking
layer_3_mode	System-to-switch trunking
balance_mode	System-to-switch load balancing enabled/disabled
timeout_mode <time>	Failover mode; time parameter in milliseconds
hub_mode	Failover mode; Only receives on the ACTIVE port. Use when connecting failover ports in the same collision domain, like to a hub. Do not use when connecting system-to-system, or system-to-switch.

To tear down the existing RAIN configuration, the interface must first be down. See ifconfig(1M):

```
ifconfig zrl0 down
rlconfig -t
```

Removing Drivers

The program `/etc/rain/Uninstall.znb` will remove all Znyx driver files except itself. It will optionally rebuild the LynxOS kernel without the driver.

Appendix

A: RAINlink for LynxOS Man Page Reference

RAINlink includes the following files and utilities to aid in configuration and monitoring:

- rainlink
- rlconfig
- rlalarm
- rlstats

This section includes Unix-style man pages to be used as reference.

rainlink(1M)

NAME

rainlink – Activate/Deactivate RAINlink

SYNOPSIS

```
/etc/rain/rainlink [ -a | -d | -s | -u ]
```

DESCRIPTION

The *rainlink* command script initializes and reinitializes the RAINlink layer within the ZNYX NetBlaster driver. An input script located in `/etc/rain/rainlink.conf` is used as input to ZNYX's `rlconfig` application. A list of default choices is presented if the `/etc/rain/rainlink.conf` input script does not exist. The script can be used to bring up the

RAINlink and non-RAINlink interfaces. It can also be used to display the current RAINlink configuration.

OPTIONS

- a Activate RAINlink features. Creates an `/etc/rain/rainlink.conf` file if it doesn't exist from a list of default choices, brings down the existing NetBlaster interfaces, tears down the existing RAINlink configuration, builds the new RAINlink configuration, brings up the RAINlink interfaces, and finally any other NetBlaster interfaces.
- d Deactivate RAINlink features. Tears down the existing RAINlink interfaces, and brings up NetBlaster interfaces.
- s Show RAINlink configuration (same as `/etc/rain/rlconfig -s`)
- u Display usage (same as no parameters)

FILES

`/etc/rain/rainlink.conf`

SEE ALSO

`rlconfig(1M)`

rlconfig(1M)

NAME

rlconfig - Configure ZNYX Redundant Array of Independent Netports (RAIN).

SYNOPSIS

```
/etc/rain/rlconfig [ -s ] [ -t ] [ -p ppa ]  
[ -l # list_file ] < input_file
```

DESCRIPTION

The *rlconfig* application is used to configure multiple ZNYX NetBlaster ports into trunks or failover groups. With exception of the *-s* option, all interfaces affected by *rlconfig* must be in the down states. See *ifconfig(1M)* for explanation on bringing down interfaces. The application reads standard input for commands.

A trunk is a grouping of two or more ports that can do port aggregation. Port aggregation can be put in a mode where multiple clients connecting to a server can utilize the bandwidth of multiple ports seamlessly, or in a mode where a server connecting to a server can seamlessly utilize the higher bandwidth of multiple ports.

A failover group is a grouping of ports where only one port is active at any one time. The remaining ports in the failover group are in standby in case the active port goes down. Typically, a failover group would only switch from one link to another if the physical link were lost. By enabling timeout mode, if no traffic is received in the specified time interval, the active port is automatically switched to another available link. The timeout mode time interval is settable by the user in increments of milliseconds.

OPTIONS

- | | |
|--------------------|---|
| -s | Displays the current RAIN configuration. Commands are not read from standard input with this option. |
| -t | Tear down all previously configured RAIN ports. Commands are not read from standard input with this option. |
| -l [#] file_name | Creates a file containing a list of the configured ZRL devices and the available ZNB devices. The “#” is used for version identification. This option is not intended for use by the user; it is used by the <i>rainlink</i> script to configure RAIN |
| -p ppa | Displays hardware level statistics for the specified ppa. |

USAGE

Typically, a script of commands is directed into *rlconfig*, which instruct *rlconfig* to build and set modes on trunks and failover groups. Semicolons delimit commands. Spaces and new lines are ignored. Commands take the following form:

```
trunk<number> = znb<number> [ , znb<number> ... ] ;  
failover<number> = znb<number> | trunk<number> [ , znb<number> |  
trunk<number> ... ] ;
```

```
trunk<number> = mode ;
trunk<number> != mode ;

failover<number> = mode ;
failover<number> != mode ;
```

where <number> is a value between 0 and the maximum number of ports in the system. Trunks can consist of ports denoted by a comma-delimited list of znb<number>. A failover group can consist of ports or trunks. You must build a trunk prior to its use in a failover group.

Acceptable modes for trunk and failover configurations:

Mode:	Appropriate for:
ip_trunking_mode	System-to-system trunking
layer_2_mode	System-to-switch trunking
layer_3_mode	System-to-switch trunking
balance_mode	System-to-switch load balancing enabled/disabled
timeout_mode <time>	Failover mode; time parameter in milliseconds
hub_mode	Failover mode; Only receives on the ACTIVE port. Use when connecting failover ports in the same collision domain, like to a hub. Do not use when connecting system-to-system, or system-to-switch.

DISPLAYS

The “-s” option of the *rlconfig* displays the current status of the RAINlink interfaces. The display includes any configured modes for trunks or failover groups, and the status of the interfaces in the form:

```
trunk<number> <mode> = znb<number>, znb<number> ... ==> zrl<number>
```

Where number is the trunk, group, or interface identifier, and mode is one of the acceptable modes. For example:

```
trunk0<IP_TRUNKING> = znb0<UP>, znb1<UP> ==> zrl0
```

The interface status can be:

Name:	Meaning:
UP	Interface is UP
DOWN	Interface is down
ACTIVE	The interface is UP, and is the ACTIVE member of a failover group
STANDBY	The interface is UP, but is not the ACTIVE member of a failover group.

The “-p” option of `rlconfig` displays the hardware level statistics for the specifies `ppa`. The display includes current `ppa` state, fault state, link configuration and current link state. Transmit and receive statistics are also displayed.

The following are valid values for the above states:

Name:	Possible states:
State:	HS_STATE_NOT_INIT HS_STATE_STOPPED HS_STATE_RUNNING
Fault:	HS_FAULT_NONE HS_FAULT_INTERNAL HS_FAULT_EXTERNAL
Link configuration	HS_LINK_AUTO HS_LINK_TP HS_LINK_BNC HS_LINK_AUI HS_LINK_TPFD HS_LINK_TX HS_LINK_TXFD HS_LINK_T4
Current link state:	HS_LINK_DOWN HS_LINK_TP HS_LINK_BNC HS_LINK_AUI HS_LINK_TPFD HS_LINK_TX HS_LINK_TXFD HS_LINK_T4

FILES

/etc/rain/rainlink.conf

SEE ALSO

rainlink(1M)

rlalarm(1M)

NAME

`rlalarm` – RAINlink switch and event monitoring application

SYNOPSIS

```
/etc/rain/rlalarm [-d <debug_level>] [-s <dev_id> <ppa>] [-r <dev_id>] [-m <ppa>]
                    [-p <dev_id> <ppa> <ppa> ...] [-h <host_name>]
```

DESCRIPTION

rlalarm is a utility based on the RAINlink Management API (RMAPI) that allows manipulation of RAINlink devices as well as a method for event notifications of a specified device. The device can be either a RAINlink or MAC level device, and can be located locally or reachable remotely via a RAINlink server daemon (`rlsd`).

OPTIONS

- s *dev_id* *ppa* Switch ACTIVE *ppa* of device *dev_id* to the specified *ppa*.
- r *dev_id* Display RAINlink notifications from specified *dev_id*.

-m ppa	Display MAC level notifications from specified <i>ppa</i> .
-p dev_id ppa ppa ...	Set precedence of device <i>dev_id</i> on failover group to order of <i>ppa</i> list. First entry in list has highest precedence.
-d debug_level	Sets debug level (1 through 4). Larger number increases verbosity.
-h host_name	Connect to remote host.

USAGE

Rlalarm allows manipulation and monitoring of RAINlink devices. To switch the active port in a failover group (zrl0) from port0 to port1, enter:

```
rlalarm -s 0 1
```

The 0 references the RAINlink (zrl) device, the 1 references the port (or ppa, znb1) to make ACTIVE in the failover group.

The "-r" and "-m" options allow you to receive notifications of event changes at either the RAINlink level or MAC level.

The "-p" option allows a method for setting a precedence within a failover group. This allows the user to configure a failover group to always return to a specified ppa, if it becomes available later. The default operation for a RAINlink failover group is to only switch when the ACTIVE port no longer is available. For example, once you've configured a failover group of two ports,

```
failover0=znb0, znb1;
```

Use "rlalarm" to specify that whenever znb0 is available, it should be the ACTIVE port:

```
rlalarm -p 0 0 1
```

You can also connect to a remote host running the RAINlink server daemon *rls*. Provide the remote hostname or ip address as an argument to rlalarm:

```
rlalarm -h 10.0.0.1 -m 0
```

This would attempt a TCP connection to remote host with IP address 10.0.0.1, and if successful retrieve the MAC level statistics for znb0 (ppa 0).

FILES

None

rlstats(1M)

NAME

rlstats – Displays statistics for NetBlaster or RAINlink devices

SYNOPSIS

```
/etc/rain/rlstats [[-m|-s|-o|-i] <ppa>] [-d <level>] [-z]
```

DESCRIPTION

The *rlstats* application is used to display statistics and configuration information for RAINlink devices.

OPTIONS

-m ppa	Displays MAC layer statistics for specified <i>ppa</i> .
-s ppa	Turns speedometer on for specified <i>ppa</i> .
-r ppa	Sets speedometer polling rate (in ms) for specified <i>ppa</i> . Use in conjunction with <i>-s ppa</i> .
-i ppa	Retrieve hardware information for specified <i>ppa</i>
-o ppa	Display speedometer statistics for the specified <i>ppa</i>
-d debug_level	Sets debug level (1 through 4). Larger number increases verbosity.
-z	Displays RAINlink configuration

USAGE

Use `rlstats` to display statistics for a particular port or RAINlink configuration.

FILES

None

Other Items

- The current release of the driver is v2.3.7. To determine the version of driver installed, use the `rlstats` command,

```
/etc/rain/rlstats -i 0
```

- By default, each network interfaces will auto-negotiate its link speed, and full or half-duplex operation. If a fixed setting is desired, edit the `media` variable in the `zxe_dev_install` program to specify the desired setting. The possible settings are listed in the program. This program is located in the directory:

```
/sys/drivers.cpci_drm/if_zxe
```

The following values set the mode for all interfaces:

Value	Mode
0	Auto-Negotiation
1	10BaseT, Half Duplex
4	10BaseT, Full Duplex
5	100BaseTx, Half Duplex
6	100BaseTx, Full Duplex

Where to go for Technical Support

The best way to obtain technical support is to email support@znyx.com.
Please include:

1. Exact model number of ZNYX hardware.
2. Exact version of the operating system.
3. A very detailed description of your problem.
4. Attaching console output in the vicinity of the problem is a big plus.

Resource	Address
Telephone	(510) 249-0800
Toll-Free	(800) 724-0911 (USA Only)
FAX	(510) 656-2460
Website	www.znyx.com
E-mail	support@znyx.com



ZNYX Corporation
48421 Millmont Drive
Fremont, CA 94538 USA

Document # DC012002
March 1, 2001

© 2000-2001 ZNYX Networks. All rights reserved worldwide. All information in this document is subject to change without prior notice. ZNYX and NetBlaster are trademarks or registered trademarks of ZNYX Networks in the United States and/or other countries. All other marks, trademarks or service marks are the property of their respective owners.

ZNYX Networks may have patents, pending patent applications, trademarks, copyrights or other intellectual property rights covered in the subject matter of this document. By furnishing this document, ZNYX Networks does not license nor waive its license to those intellectual property rights except as expressly provided in a written license agreement from ZNYX Networks. Information in this document is subject to change without prior notice.